

Optimal Oblivious Parallel RAM

Wei-Kai Lin


Joint work with

Gilad Asharov, Ilan Komargodski, Enoch Peserico, Elaine Shi










Access pattern
leaks data



Frequency,
Correlation









Oblivious algorithms:
Locations “indep of”
(secret) data




[Goldreich-Ostrovsky87,96]



 Accessed
locations

Infinite algos / data structs?
Yes, **Oblivious RAM**









Operation: read or write



Accessed
locations




Overhead:
Num. accesses per operation




ORAM Overhead, Standard Setting

$O(\sqrt{n})$



[G87]

$O(\log^3 n)$



[GO96]
[DMN11]
[SCSL11]

$O(\log^2 n)$



[GM11]
[SvDSF+13]
[WCS15]

$O(\log^2 n / \log \log n)$



[KLO12]


$O(\log n \cdot \log \log n)$



[PPRY18]

- n : size of the memory
- Word size: $\log n$ bits
- Client's memory size $O(1)$ words
- Assume existence of cryptographic pseudorandom function

$O(\log n)$




[AKLNPS20]
"OptORAMa"

$\Omega(\log n)$






[GO96]
[LN18]


Problem solved




Parallel RAM?
Oblivious PRAM

 Accessed locations





Batch of operations





Oblivious Parallel RAM (OPRAM)



Accessed locations

Same num of CPUs

Overhead:
Num. parallel steps
per batch of operations



Main question: OPRAM overhead?

- n : size of the memory
- Word size: $\log n$ bits
- Client's memory size $O(1)$ words
- Assume existence of cryptographic pseudorandom function

$O(\log^3 n)$



[BCP16]

$O(\log^2 n)$



[KN16]
[CLT16]

$O(\log^2 n / \log \log n)$



[CS17]
[CCS17]
[CGLS17]

$\log n?$



Batched,
 $< \log n?$




[GO96]
[LN18]

Main result:
Optimal OPRAM,
 $\Theta(\log n)$ overhead

- n : size of the memory
- Word size: $\log n$ bits
- Client's memory size $O(1)$ words
- Assume existence of cryptographic pseudorandom function

$O(\log^3 n)$



[BCP16]

$O(\log^2 n)$




[KN16]
[CLT16]

$O(\log^2 n / \log \log n)$



[CS17]
[CCS17]
[CGLS17]

$O(\log n)$



$\Omega(\log n)$




Num CPU
< $n^{0.99}$

$\Omega(\log n)$






[GO96]
[LN18]



Hierarchical Framework

[Goldreich-Ostrovsky '87,96]

Level i , capacity 2^i




Put ● into Level 0

Rebuild (hash)
level i per 2^i operations

Oblivious Hash Table:

- “Build”: store 2^i balls
- “Lookup” for key k :
find and output k and
associated ball
- Build + Lookup is oblivious:
not reveal balls even
adversary sees accesses




Optimal ORAM (sequential)



[PPRY18]
"PanORAMa"

[AKLNPS20]
"OptORAMa"

Level i , capacity 2^i



Rebuild (hash)
level i per 2^i requests



Oblivious Hash Table


- Build: **linear** time
- Lookup: **const** time

Linear time building blocks

"Tight compaction":
Oblivious sorting balls using 1-bit keys

"Intersperse":
Mixing 2 lists uniformly at random

"Toss balls into bins,
and then report loads of bins"



Oblivious **Parallel** RAM


Also need:
 $O(\log n)$ parallel time

- Oblivious Hash Table
- Build: **linear work**
 - Lookup: const work

Level i , capacity 2^i
0



1, H_1




Yes! [AKLPS20]*

Linear work building blocks

“Tight compaction”:
Oblivious sorting balls using 1-bit keys


2, H_2



Yes! [This work]

“Intersperse”:
Mixing 2 lists uniformly at random


3, H_3



Yes! [This work]




“Toss balls into bins,
and then report loads of bins”

*Related: “optimal sorting circuits” (Wednesday 4:30, Session 12B)



Oblivious **Parallel** RAM

Level i , capacity 2^i



Building blocks take:
Linear work
 $O(\log n)$ parallel time

*log n levels,
each takes
log n parallel time*

Pipeline levels [This work]

Main Technical Challenge


\Rightarrow

OPRAM:
 $\sim \log^2 n$ overhead

\Rightarrow

OPRAM:
 $\log n \cdot \text{poly log log } n$
overhead

- Oblivious Hash Table
- Build: linear work,
 $\log n \cdot \text{poly log log } n$ parallel time
 - Lookup: const work




Oblivious Hash Table

[PPRY18] [AKLNPS20]

1. "Balls into bins" hashing

Choose the bin by $H(\text{addr})$



Poly log bin-size \rightarrow negligible overflow prob.

2. "Cuckoo" hashing within each bin

Build cuckoo hashing:
linear work,

\times $O(\log n \cdot \text{poly log log } n)$
parallel time


Cuckoo:
Const-time
lookup

Level

0



1, H_1



$H_1(5)$

2, H_2




$H_2(5)$

3, H_3




$H_3(5)$




Level
0




1, H_1



2, H_2




3, H_3



“Balls into bins” hashing

Choose the bin by addr



“Short Hash Table”



Wanted Hash Table

- Build: linear work,
 $\log n$ parallel time
- Lookup: const work

A. Random shuffled input

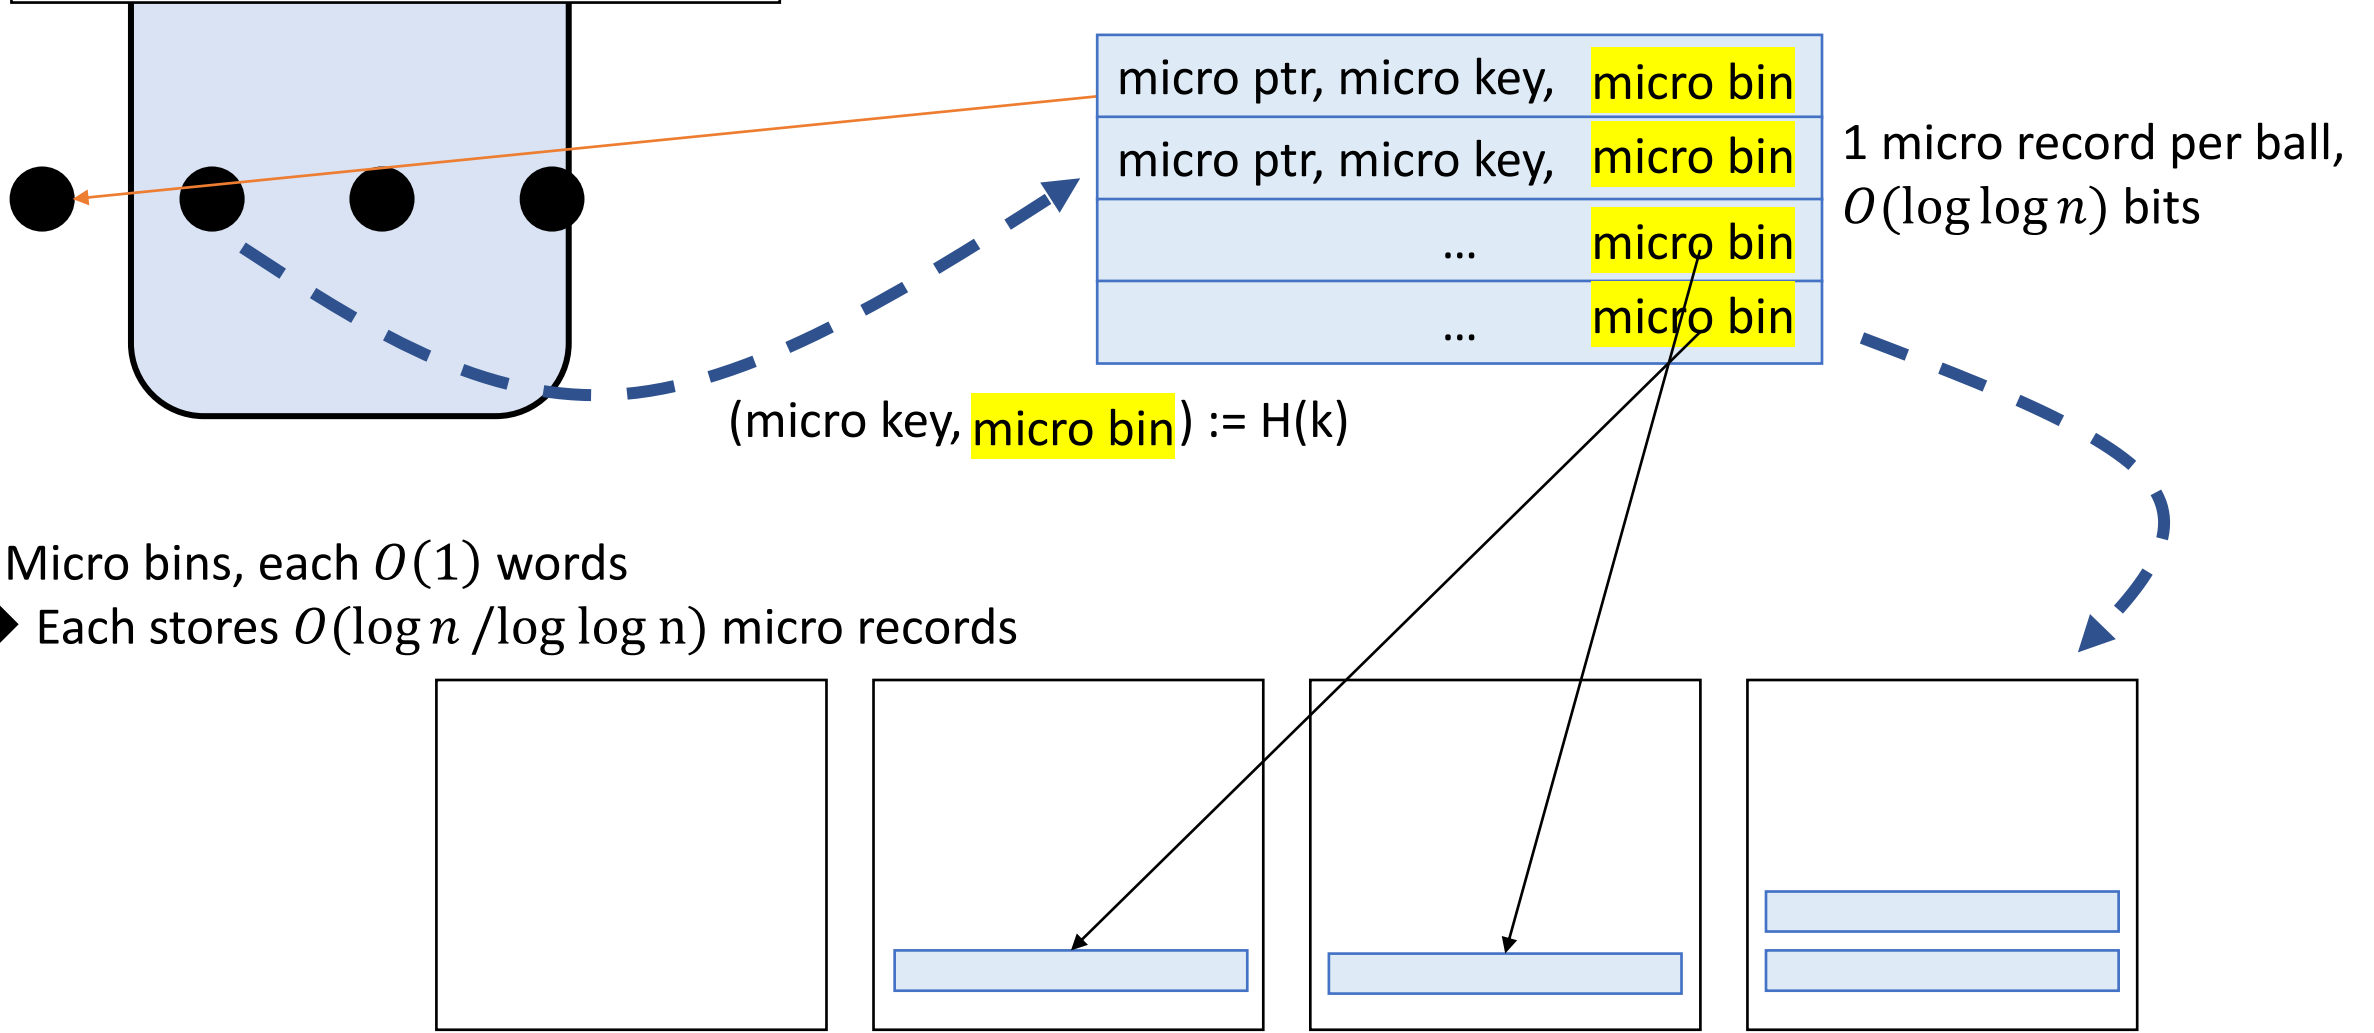
B. “Amortized” over many bins

Build many bins


Lookup $\log n$ bins

- Word size: $\log n$ bits
- Ball: (key, value) pair, $\log n$ bits
- Bin size: $s = \text{poly log } n$ words
- Num balls: $O(s)$

A. Random shuffled input



- Word size: $\log n$ bits
- Ball: (key, value) pair, $\log n$ bits
- Bin size: $s = \text{poly log } n$ words
- Num balls: $O(s)$

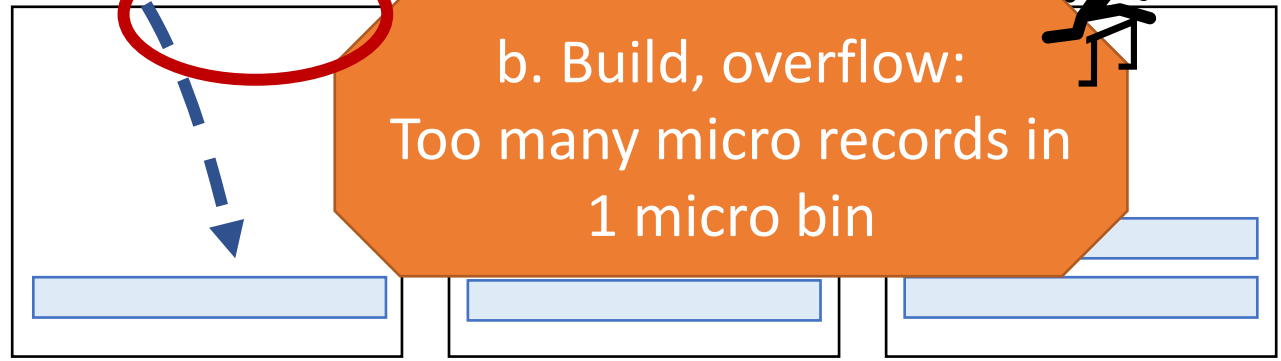


a. Build: Need time to pick **unique micro keys**

micro ptr **micro key**, micro bin


1 micro record per ball, $O(\log \log n)$ bits
 (micro key, micro bin) := $H(k)$

b. Build, overflow: Too many micro records in 1 micro bin



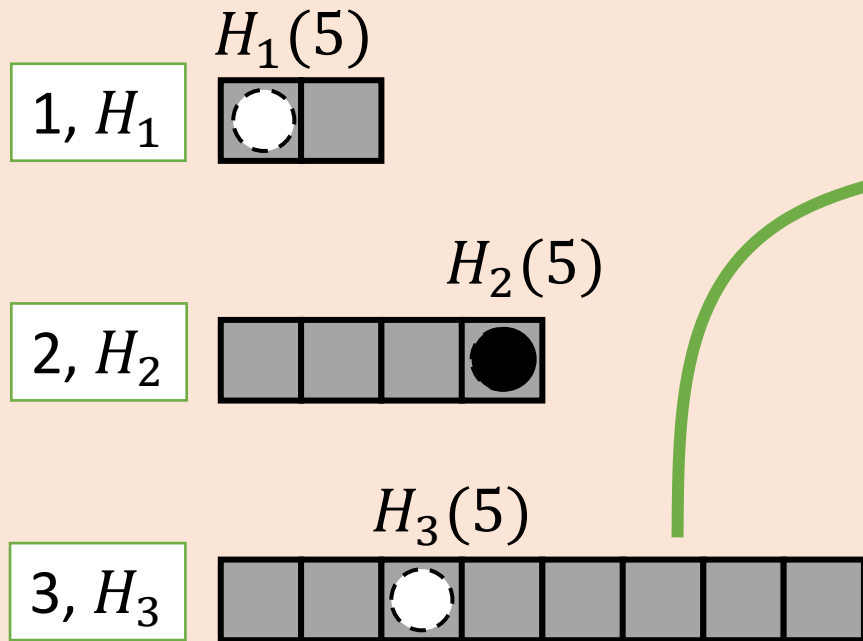
s Micro bins, each $O(1)$ words
 → Each stores $O(\log n / \log \log n)$ micro records

c. Lookup, false positive: key' not in this index hits same micro key

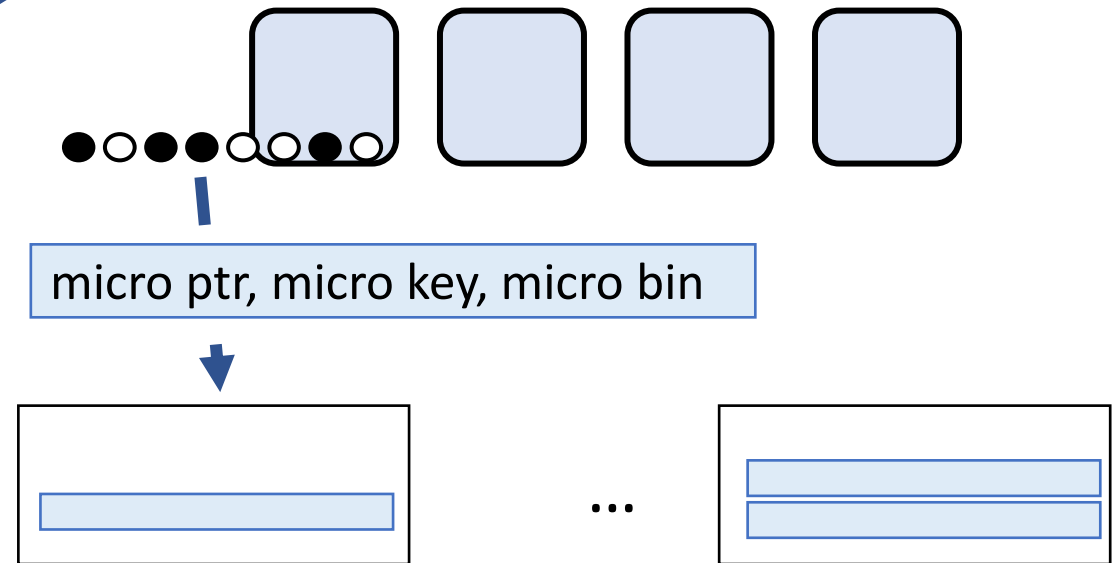


💡 - B. "Amortized" over many bins

- Build:
 $> poly \log n$ instances of ShortHTs (in parallel)
- Lookup:
 $\log n$ instances of ShortHTs (sequential)
- Total number of instances $< n$
➔ share CPUs, overflow/false positive space

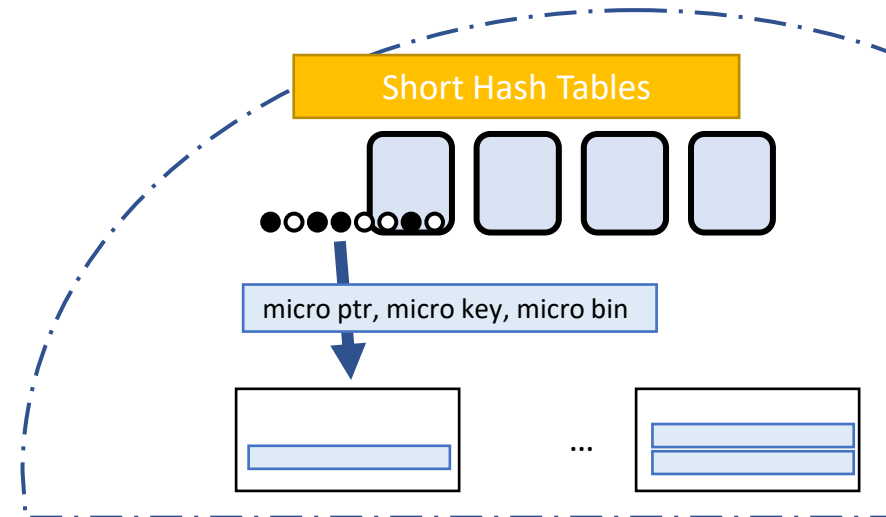



Short Hash Tables



Omitted Details (need oblivious and parallel):

- Pipeline $\log n$ levels of hash tables into Oblivious PRAM
- Building blocks
- Lookup Short Hash Table concurrently



- Extend lower bounds of [GO96] and [LN18] to parallel

Thank you!